

Shape Constraints in Economics and Operations Research

Andrew L. Johnson and Daniel R. Jiang

Texas A&M University and Osaka University
University of Pittsburgh

Abstract. Assumptions motivated by either logical or existing theory can be imposed during model estimation to restrict the feasible region of the parameters. The restrictions, implemented as shape constraints, may not provide any benefits in an asymptotic analysis, but will improve the estimator's finite sample performance. This paper briefly reviews an illustrative set of research on shape constrained estimation in the economics and operations research literature. We highlight the methodological innovations and applications, with a particular emphasis on utility functions, production economics, and sequential decision making applications.

Key words and phrases: Shape Constraints, Multi-variate Convex Regression, Nonparametric Regression, Production Economics, Consumer Preferences, Revealed Preferences, Approximate Dynamic Programming, Reinforcement Learning.

1. INTRODUCTION

This paper builds on prior surveys of shape constrained estimation in the economics literature by [Matzkin \(1994\)](#), [Yatchew \(2003\)](#), and [Chetverikov et al. \(2018\)](#), and also surveys the operations research literature for the first time. Although length limitations prohibit a comprehensive survey, we describe the important central themes and identify the recent advances and applications and active research directions in the literature.

One of the first papers published in the economics literature is [Hildreth \(1954\)](#) who estimated the relationship between corn output and nitrogen fertilizer restricting output to be a function of fertilizer that was monotonically increasing and concave. Later, [Brunk \(1970\)](#) studied isotonic functions, or imposing a monotonicity constraint on the function, and showed for isotonic regression with a single regressor the consistency and the rate of convergence. [Barlow et al. \(1972\)](#)

Andrew Johnson is an Associate Professor in the Department of Industrial and Systems Engineering, Texas A&M University, College Station, Texas 77840, USA, and holds an appointment as a Visiting Associate Professor in School of Information Science and Technology, Osaka University, Suita 565-0871, Japan. (e-mail: ajohnson@tamu.edu). Daniel Jiang is an Assistant Professor in the Department of Industrial Engineering, University of Pittsburgh, Pittsburgh, PA 15261, USA. (e-mail: drjiang@pitt.edu).

were the first to develop and organize fundamental results in order restricted inference for isotonic regression. [Robertson et al. \(1988\)](#) summarized one of the earliest conferences on shape constrained functional estimation. [Mammen \(1991\)](#) considered the two-step estimation of a smooth monotone function in a single regressor setting and analyzed interchanging the isotonization step and the smoothing step. [Hall and Huang \(2001\)](#) considered the estimation of a monotonic and convex/concave function with a single regressor using kernel weighting methods with an additional weigh vector to assure the function satisfied the shape constraints. [Villalobos and Wahba \(1987\)](#), who considered the estimation of a smooth function with a set of linear inequality constraints, provided a characterization which could be solved using a nonlinear programming algorithm.

Our survey is divide into a section on regression-based estimation and a section on sequential decision making. Within the regression-based estimator section, we divide the relevant research into three subsection: 1) The Nonparametric Least Squares Estimator; 2) Approximate Global Shape Restrictions, Ensemble Methods, and Local Weighting; and 3) Single Index Models and Alternative Assumptions to Global Convexity. In each subsection we describe both a set of estimators and an application that illustrates the estimators. For subsection 1) the application is utility function estimation, whereas in subsections 2) and 3) the application is production function estimation. In the sequential decision making section, the subsections are 1) Convexity of the Value Function, 2) Monotonicity and 3) Policy Structure. Several examples are included throughout the section. We refer readers interested in other relevant topics and applications of shape constraint estimation, e.g., option pricing or stochastic approximation to see [Art-Sahalia and Duarte \(2003\)](#) and [Kushner and Yin \(2003\)](#), respectively. For testing, we suggest [Chernozhukov et al. \(2015\)](#) and [Chetverikov et al. \(2018\)](#) and the references therein.

2. REGRESSION-BASED ESTIMATORS

Regression models allow the investigation of observational data to identify correlations among variables to provide basic empirical evidence of the relationship between variables. Here we will review several regression based methods for estimating these relationship while imposing additional information which can be stated as shape constraints on the regression models. We will begin in Section [2.1](#) with the widely used nonparametric least squares estimator (LSE). Estimating the nonparametric LSE with shape constraints results in functional estimates that satisfy a set of global axiomatic properties. We will describe the application of using Afriat inequalities to test the generalized axioms of revealed preferences. Then in Section [2.2](#) we will consider approximations to nonparametric estimator with globally axiomatic properties. These estimators will approximate either the axioms or the function by only imposing weaker versions of the axiom or by using parametric approximations. In some cases these approximations can be improved by ensemble methods or local weighting both of which will be discussed. We will discuss applications related to production economics where these estimators are used. Finally in Section [2.3](#), we will review single index models and alternative shape restrictions. Single index models assure predictor variables are aggregated to maintain properties such as linearity or convexity while allowing the relationship between the dependent variable and the aggregate predictor

to remain general. These models can be used to implement more complicated economic models of production where convexity is maintained between certain variables while more complicated relationships such as the S-shape can be imposed between other variables. We will illustrate these models with extensions developed in production economics applications.¹

2.1 The Nonparametric Least Squares Estimator

One of the most widely used shape constrained estimators is the nonparametric LSE of a multi-variate convex regression function (Groeneboom et al., 2001; Kuosmanen, 2008; Seijo and Sen, 2011). Initial work on multi-variate convex regression functions includes Matzkin (1991), Banker and Maindiratta (1992), and Allon et al. (2007), who considered maximum likelihood estimators. The consistency of the first two estimators is shown in Matzkin (1991) and Sarath and Maindiratta (1997), respectively. However, both estimators have had little practical implementation because of the computational complexity of the associated optimization problems. Allon et al. (2007) described some of these complexities and used the concept of entropic distance to develop a maximum likelihood estimator that can be stated as a convex programming problem and thus can be solved for data sets of up to 400 observations which was considered a large instance. Many of the computational strategies for the nonparametric LSE described below are directly applicable to the Allon estimator, but they have not been pursued in the literature. Alternatively, Beresteanu et al. (2007) considered a sieve estimator with a least squares loss function, and evaluated the metric entropy of the space of shape-restricted functions.

Kuosmanen (2008) proposed the characterization of the nonparametric LSE of a multi-variate convex regression function. Unlike the previous estimators, this characterization relaxes assumptions on the distribution of the error term and the need for turning parameters. To define the model, consider the set of observations $\{(\mathbf{X}_i, Y_i) : i = 1, 2, \dots, n\}$ and a nonparametric shape restricted regression satisfying

$$(2.1) \quad Y_i = f(\mathbf{X}_i) + \epsilon_i.$$

Here, $\mathbf{X}_i \in \mathbb{R}^d$ is an observed vector of predictors where $d \geq 1$. The noise term, ϵ_i satisfy $\mathbf{E}(\epsilon_i | \mathbf{X}) = 0$, and the real-valued regression function f is unknown but obeys certain known restrictions. In the work described to this point, the known restrictions were convexity and in some cases also monotonicity. Thus, let \mathcal{F} denote the class of all regression functions satisfying a particular set of restrictions. Letting $\boldsymbol{\theta}^* = (f(\mathbf{X}_1), f(\mathbf{X}_2), \dots, f(\mathbf{X}_n))$, $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$ and $\boldsymbol{\epsilon} = (\epsilon_1, \epsilon_2, \dots, \epsilon_n)$, rewrite the model (2.1) as

$$\mathbf{Y} = \boldsymbol{\theta}^* + \boldsymbol{\epsilon}.$$

Subject to the constraints imposed by the properties of \mathcal{F} , the estimation problem can be translated to constraints on $\boldsymbol{\theta}^*$ of the form $\boldsymbol{\theta}^* \in \mathcal{C}$. The set \mathcal{C} contains all possible \mathbf{Y} that can be generated by f , the set of functions in the shape restriction

¹Throughout this section we will use the mathematical notation common in this literature which is to indicate matrices and vectors by using a bold font and indicate scalar variables as non-bold font. Other notation is introduced as it is used.

class, given the observed matrix \mathbf{X} . Specifically

$$\mathcal{C} = \{(f(\mathbf{X}_1), f(\mathbf{X}_2), \dots, f(\mathbf{X}_n)) \in \mathbb{R}^n : f \in \mathcal{F}\}$$

is a closed subset of \mathbb{R}^n . Now consider an estimator of $\boldsymbol{\theta}^*$. Define the LSE estimator $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}^*$ for a shape-restricted regression as the projection of \mathbf{Y} onto the set $\boldsymbol{\theta}^*$,

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}^* \in \mathcal{C}} \|\mathbf{Y} - \boldsymbol{\theta}\|^2,$$

where $\|\cdot\|$ denotes the Euclidean norm in \mathbb{R}^n . Generally, because \mathcal{C} is a closed convex set, $\hat{\boldsymbol{\theta}} \in \mathcal{C}$ is unique and it can be characterized by:

$$\langle \mathbf{Y} - \hat{\boldsymbol{\theta}}, \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \rangle \leq 0, \text{ for all } \boldsymbol{\theta} \in \mathcal{C},$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product in \mathbb{R}^n . Consider the specific case of convex regression when the restriction on \mathcal{F} is convexity. As discussed in [Seijo and Sen \(2011\)](#), both primal and dual characterizations are possible for the LSE's constraints. Pursuing the dual characterization leads to the positive semidefinite quadratic program:

$$(2.2) \quad \begin{aligned} & \underset{\boldsymbol{\theta}, \boldsymbol{\beta}}{\text{minimize}} && \sum_{k=1}^n (Y_k - \theta_k)^2 \\ & \text{subject to} && \langle \boldsymbol{\beta}_k, \mathbf{X}_j - \mathbf{X}_k \rangle \leq \theta^j - \theta^k \quad \forall k, j = 1, 2, \dots, n. \end{aligned}$$

Here, $\boldsymbol{\beta}_k$ is the estimated subgradient at the point \mathbf{X}_k . In this paper, we refer to the constraints in (2.2) as the *Afriat inequalities* ([Afriat, 1967, 1972](#)).

It has been noted that LSE achieves limited performance near the boundaries of the domain ([Seijo and Sen, 2011](#); [Lim and Glynn, 2012](#)). [Lim \(2014\)](#) proposed a refinement to include restricting the domain of \mathcal{F} and bounding the subgradients of f , resulting in the set

$$\begin{aligned} \mathcal{F} = \{ & f : [0, 1]^d \rightarrow \mathbb{R} : f \text{ is convex and } |\xi^j| \leq C \text{ for } 1 \leq j \leq d, \\ & \text{where } (\xi^1, \xi^2, \dots, \xi^d) \in \partial f(\mathbf{x}) \text{ for } \mathbf{x} \in (0, 1)^d \}, \end{aligned}$$

where $\partial f(\mathbf{x})$ is the subdifferential of f at \mathbf{x} . This means that additional constraints on the estimated subgradients can be added to (2.2):

$$(2.3) \quad |\boldsymbol{\beta}_k| \leq C \quad \forall k = 1, 2, \dots, n.$$

Let \hat{f} be an estimate of the function f computed using (2.2) augmented with (2.3). [Lim \(2014\)](#) establishes a rate of convergence result: for any $\varepsilon > 0$, there exists a constant c , depending on C and ε , that satisfies

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left[\left(\frac{1}{n} \sum_{j=1}^n (\hat{f}(\mathbf{X}_j) - f(\mathbf{X}_j))^2 \right)^{1/2} \geq r_n c \right] < \varepsilon,$$

where

$$r_n = \begin{cases} n^{-2/(4+d)}, & \text{if } d < 4. \\ (\log n)^{1/2} n^{-1/4}, & \text{if } d = 4. \\ n^{-1/d}, & \text{if } d \geq 4. \end{cases}$$

Juditsky and Nemirovski (2002) obtain a similar but more general result by estimating the distance from an unknown signal in a white-noise model to the convex cones of the positive/monotone/convex function. The authors showed that when the function belongs to a Hölder class, the risk of estimating the L_r -distance, $1 \leq r < \infty$, from the signal to the cone is the same, up to a logarithmic factor, as estimating the signal itself. The same risk result holds for testing if an unknown signal is positive, monotonic, and/or convex. See Guntuboyina and Sen (2013, 2015) and Chatterjee et al. (2015) for other results related to the rate of convergence.

2.1.1 Computation Wu (1982) and Dykstra (1983) published the initial work on computational issues for the least squares estimator subject to convexity with a single regressor. Goldman and Ruud (1993) considered the multi-variate problem and recognized that Dykstra’s insight could be generalized from searching the intersection of a finite number of convex cones to search over the intersection of a finite number of convex sets. More recent operations research literature investigated LSE’s computational issues. For example, Alizadeh (2006) discussed the connection between shape constrained functional estimators and the corresponding semidefinite constraints on the programming problems used to calculate LSE’s parameters.

Lee et al. (2013) noted that while (2.2) is in the class of convex optimization and is thus polynomial-time solvable, the number of constraints grows quadratically with the size of the data set and can thus present practical challenges even for a few hundred observations. The percentage of binding constraints at the optimal solution is only 0.5 - 1% for typical problems in economic applications. Therefore the authors proposed an estimation procedure which initially included only the constraints for observations that are close to one another in terms of Euclidean distance $\|\mathbf{X}_i - \mathbf{X}_j\|^2 < C$ where C is prespecified parameter. Notating the set of constraints implied by the closeness criteria as V , the constraints are

$$(2.4) \quad \langle \boldsymbol{\beta}_k, \mathbf{X}_j - \mathbf{X}_k \rangle \leq \theta^j - \theta^k \quad \forall k, j \in V, \boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \dots, \boldsymbol{\beta}_n \in \mathbb{R}^d, \boldsymbol{\theta} \in \mathbb{R}^n.$$

To solve a series of smaller optimization problems, Lee et al. (2013) proposed Algorithm 1 below, which adds back violated constraints iteratively until satisfying the set of all Afriat constraints. The algorithm is demonstrated for data sets of up to 1,800 observations.

Alternatively, using a least absolute deviation loss function, which can be formulated as an easily solvable linear program, Luo and Lim (2016) showed that the estimator converges almost surely to the true function as n increases to infinity. Mazumder et al. (2015) presented an Alternating Direction Method for Multipliers (ADMM) algorithm for solving the nonparametric LSE of a multi-variate convex regression function. Although their algorithm calculates estimates for data sets of 3,500 observations in less than one hour, their proposed ADMM algorithm has not been proven to converge for all data sets. Specifically, the ADMM algorithm requires dividing the variables into groups and the algorithm alternates between fixing all but one of the groups’ variables to their current best values and optimizing only in terms of the selected group of variables. The ADMM algorithm has been proven to converge for two groups of variables; however, Mazumder et al. (2015) uses three groups of variables and the converge properties of ADMM with three groups is still an open question (Bertsekas, 1999).

Algorithm 1 Iterative algorithm proposed by Lee et al. (2013)

- (1) Let $t = 0$ and let V be a subset of the ordered pairs of observations.
 - (2) Solve (2.2) replacing the Afriat inequalities with the (2.4) to find an initial solution, $(\boldsymbol{\theta}^{(0)}, \boldsymbol{\beta}^{(0)})$.
 - (3) Do until $(\boldsymbol{\theta}^{(t)}, \boldsymbol{\beta}^{(t)})$ satisfies all the Afriat inequalities:
 - (a) Select a subset of the concavity constraints that $(\boldsymbol{\theta}^{(t)}, \boldsymbol{\beta}^{(t)})$ violates and let $V(t)$ be the corresponding observation pairs.
 - (b) Set $V = V \cup V^{(t)}$.
 - (c) Solve (2.2) replacing the Afriat inequalities with the (2.4) and using V from (b) to obtain solution $(\boldsymbol{\theta}^{(t+1)}, \boldsymbol{\beta}^{(t+1)})$.
 - (d) Set $t = t + 1$.
-

2.1.2 Application: Consumer Preferences Standard approaches to analyze the demand-side of an economy typically assume parametric functions for consumer preferences and demand functions to develop models which are then fitted to the observed data. An alternative approach called revealed preferences, proposed by Samuelson (1938), used nonparametric methods, thus avoiding parametric assumptions for both preferences and demand. In this section, we present the concepts of rational preferences and the generalized axiom of revealed preferences (GARP). We explain how the Afriat inequalities, the shape constraints mentioned above, can be used to test GARP. We discuss several applications of shape constraints in the consumer preference literature. We end by summarizing the literature that describes how to estimate a demand functional while imposing shape constraints that imply consumer rationality.

Following the notation of Cherchye et al. (2007), under the assumption of a set of T price-quantity pairs, denote the vectors of the prices and quantities associated with observation t by \mathbf{p}_t and \mathbf{q}_t , respectively, where $\mathbf{p}_t \in \mathbb{R}_{++}^N$ and $\mathbf{q}_t \in \mathbb{R}_+^N$. The observations describe how the price \mathbf{p}_t changes over time and a rational consumer should adjust their consumption \mathbf{q}_t accordingly.²

DEFINITION 1 (Samuelson (1950)). *Let $\mathcal{S} = \{(\mathbf{p}_t, \mathbf{q}_t)\}_{t=1}^T$ be a set of observations. A utility function U provides a rationalization of \mathcal{S} if for each observation t , we have $U(\mathbf{q}_t) \geq U(\mathbf{q})$ for all \mathbf{q} with $\mathbf{p}_t^\top \mathbf{q} \leq \mathbf{p}_t^\top \mathbf{q}_t$.*

Note $\mathbf{p}_t^\top \mathbf{q}_t$ is the total that a consumer pays for quantity \mathbf{q}_t . The definition states that for a utility function to be rational, if the consumer is willing to pay more, $\mathbf{p}_t^\top \mathbf{q} \leq \mathbf{p}_t^\top \mathbf{q}_t$, then the utility should be larger, $U(\mathbf{q}_t) \geq U(\mathbf{q})$. And if the data set, \mathcal{S} can be rationalized, then the consumer behaved rationally during the observed time period.

Define the property, *local nonsatiation*, to mean that for any bundle of goods there is always another bundle of goods arbitrarily close that is preferable, i.e., a consumer always prefers more of or less of an item and not a specific amount. Note

²We use \mathbb{R}_+^N to refer to the N dimensional positive orthant including the origin and \mathbb{R}_{++}^N is the strictly positive orthant.

strong monotonicity implies local nonsatiation but not vice versa. The revealed preferences literature has defined a locally nonsatiated utility function that provides a rationalization of the set of observations S if and only if the data satisfy the *Generalized Axiom of Revealed Preferences (GARP)*.

In order to define GARP, we first introduce the notions of *strictly directly revealed preferred* and *revealed preferred*. Let $\mathcal{S} = \{(\mathbf{p}_t, \mathbf{q}_t)\}_{t=1}^T$ be a set of observations. The bundle of quantities \mathbf{q}_t is *strictly directly revealed preferred* to \mathbf{q} if $\mathbf{p}_t^\top \mathbf{q}_t > \mathbf{p}_t^\top \mathbf{q}$. Moreover, \mathbf{q}_t is *revealed preferred* to \mathbf{q} if

$$\mathbf{p}_t^\top \mathbf{q}_t \geq \mathbf{p}_t^\top \mathbf{q}_u, \mathbf{p}_u^\top \mathbf{q}_u \geq \mathbf{p}_u^\top \mathbf{q}_v, \dots, \mathbf{p}_z^\top \mathbf{q}_z \geq \mathbf{p}_z^\top \mathbf{q},$$

for some sequence (t, u, \dots, z) . The set of observations S satisfies GARP if for any t and s , the bundle \mathbf{q}_t being revealed preferred to \mathbf{q}_s implies that \mathbf{q}_s is not strictly directly revealed preferred to \mathbf{q}_t . With these definitions, we can formalize the relationship between GARP and the existence of a utility function that can rationalize the data set S .

THEOREM 2.1 (Afriat (1967); Diewert (1973); Varian (1982)). *For a set $\mathcal{S} = \{(\mathbf{p}_t, \mathbf{q}_t)\}_{t=1}^T$ of observations of price-quantity pairs, the following statements are equivalent:*

- (1) *There exists a utility function U , satisfying local nonsatiation, that provides a rationalization of \mathcal{S} .*
- (2) *The set \mathcal{S} satisfies GARP, as defined above.*
- (3) *There exist $U_1, \lambda_1, U_2, \lambda_2, \dots, U_T, \lambda_T \in \mathbb{R}_{++}^N$ such that the Afriat inequalities hold: for all $t, r \in \{1, 2, \dots, T\}$,*

$$U_r - U_t \leq \lambda_t \mathbf{p}_t^\top (\mathbf{q}_r - \mathbf{q}_t).$$

- (4) *There exists a continuous monotonically increasing and concave utility function U that satisfies local nonsatiation and provides a rationalization of \mathcal{S} .*

Condition (2) implies that GARP is necessary and sufficient for rationalization of the data. While Condition (3) states that satisfying the Afriat inequalities is equivalent to satisfying GARP. Further, Condition (3) provides a test for GARP; specifically, if the Afriat inequalities are satisfied, then the set of observations \mathcal{S} satisfy GARP. The utility function referenced in (1) is defined as the lower envelope of a set of hyperplanes, specifically

$$U(x) = \min_t U_t + \lambda_t \mathbf{p}_t^\top (\mathbf{q} - \mathbf{q}_t).$$

Brown and Matzkin (1996) extended Afriat results to other important economic models; specifically, they identified restrictions on prices, incomes, and endowments for general equilibrium models.

GARP and the associated estimators are the workhorses for the nonparametric analysis of demand functions. Numerous extensions have been developed, such as Matzkin (1991), who extended revealed preference analysis to nonlinear budget sets while still imposing concave utility functions. Blundell et al. (2003) studied both observations and experimental data and developed methods to detect revealed preference violations while considering the evolution of consumer preferences over time modeled as an expansion path. Blundell et al. (2008), who

considered a case with a small number of market prices and a large number of consumers within each market, used consumer level data to estimate the bounds on consumers' responses to new relative prices.

Note that if the GARP test is applied across households, it is often rejected. This simply implies that all households do not have the same utility function and thus have different demand functions. Often economists are interested in estimating a consumer demand equation, $D(\mathbf{Z}_i) + b_i$. Here, the demand function, $D(\cdot)$ is typically a function of the prices, total expenditures and observable characteristics of household, \mathbf{Z}_i , plus an individual-specific term b_i . [Lewbel \(2001\)](#) recognized that GARP should be applied to an individual's demand function and not the aggregate demand function that is part of the consumer demand equation. Thus, [Lewbel \(2001\)](#) listed the conditions for the demand function from the consumer demand equation to satisfy utility function rationalization.

The analysis of heterogeneous demand using nonparametric shape constrained methods has recently renewed in popularity. Examples include [Dette et al. \(2016\)](#), who proposed a test for downward sloping demand curves, and [Hausman and Newey \(2016\)](#), who considered unobserved individual heterogeneity in utility functions and focused on recovering the relationship between demand and price for a continuous consumption good. As with many nonparametric shape constrained estimators, [Hausman and Newey \(2016\)](#) did not achieve point-identification rather they obtained bounds on the average welfare effect. [Bhattacharya \(2016\)](#) studied discrete demand with general heterogeneity in customers and derived the average consumer surplus.

2.2 Approximations to Global Shape Restrictions, Ensemble Methods, and Local Weighting

There are two potential limitations of the nonparametric shape constrained estimators described above: 1) the estimators are computationally expensive and often difficult to implement, and 2) they tend to over-fit the observed sample. Below we review alternatives to address these two potential limitations.

2.2.1 Approximation to Global Shape Restrictions [Gallant and Golub \(1984\)](#) proposed imposing shape restrictions, such as convexity or quasi-convexity, for a Fourier Functional Form ([Gallant, 1981](#)). The Fourier Functional Form is an early example of a sieve estimator, which approximates a sequence of parameter spaces in which the parameter space increases as the number of observations increases. In the limit, as the number of observations $n \rightarrow \infty$, the approximating parameter space is dense in the original space. [Diewert and Wales \(1987\)](#) argued that economic theory implies *global* curvature properties which are not guaranteed by Gallant's estimator. [Diewert and Wales \(1987\)](#) identified two functional forms, the Generalized McFadden and the Generalized Barnett cost functions, which, when augmented with restrictions on the matrix of second order partial derivatives, assured globally convexity and still maintained flexibility ([Diewert, 1974](#)).

Similar to [Gallant and Golub \(1984\)](#), methods to impose some shape restrictions on estimators without achieving global convexity have been introduced. For example, [Ryan and Wales \(2000\)](#) considered imposing shape restrictions only at one point in the domain of the function. The authors argued that in many cases it is sufficient to result in estimates that are globally concave; they discussed

examples of Translog and Generalized Leontief cost functions with a single regressor estimated with 25 observations that produce globally concave estimates. Alternatively, considering nonparametric shape constrained estimation, [Du et al. \(2013\)](#) proposed imposing coordinate-wise concavity which, like [Ryan and Wales \(2000\)](#), can be thought of as an approximation to global concavity. However, for more complicated models, we found very little empirical evidence that imposing concavity at a point or coordinate-wise will result in globally concave estimates or something even close.

Other models that imposed approximations to global shape restrictions include [Pya and Wood \(2015\)](#), who considered a generalized additive model under first and second order shape constraints. The estimator, which is an extension of P-splines, facilitates efficient estimation of the smoothing parameters as part of the model estimation. The authors also developed algorithms to calculate simulation-free approximate Bayesian confidence intervals for the smooth components. Similarly, [Wu and Sickles \(2017\)](#) proposed a semiparametric estimator which uses penalized splines and an integral transformation to impose monotonicity and curvature constraints. The estimator is consistent and the authors derived the asymptotic variance. Alternatively, [Chen and Samworth \(2016\)](#) considered a slightly more general model

$$Y_i = f^I(\mathbf{X}_i) = f_1(\boldsymbol{\theta}_1^\top \mathbf{X}_i) + \cdots + f_m(\boldsymbol{\theta}_m^\top \mathbf{X}_i) + c.$$

where the value of $m \in \mathbb{N}$ is assumed known, and Y_i is the response variable and follows an exponential family distribution. The variable $c \in \mathbb{R}$ is the intercept term, $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_m \in \mathbb{R}^d$ are called the projection indices, and $f_1, f_2, \dots, f_m : \mathbb{R} \rightarrow \mathbb{R}$ are called the ridge functions. [Chen and Samworth \(2016\)](#) generalizes the class of generalized additive models by allowing each function f to be a function of a linear aggregation of \mathbf{X} , specifically to the component m , $\boldsymbol{\theta}^\top \mathbf{X}_i$. The benefit of this model is that shape constraints can be imposed on each of the single-dimensional functions f , making the estimator scalable and reducing the computational difficulty. The drawback is that shape are only imposed coordinate-wise and global properties are not imposed.

2.2.2 Other Shape Constrained Estimators and Ensemble Methods [Hannah and Dunson \(2011\)](#) proposed Multi-variate Bayesian Convex Regression (MBCR), which approximated a general convex multi-variate regression function with the maximum value of a random collection of hyperplanes. Additions, removals, and changes of proposed hyperplanes are done through a Reversible Jump Markov Chain Monte Carlo (RJCMC) algorithm ([Green, 1995](#)). One of MBCR's attractive features include the block nature of its parameter updating, which causes the parameter estimate autocorrelation to drop to zero in tens of iterations in most cases. In addition, MBCR spans of all convex multi-variate functions without the need for any acceptance-rejection samplers, scales to a few thousand observations, and relaxes the homoscedastic noise assumption.

[Magnani and Boyd \(2009\)](#) proposed an iterative fitting scheme to select the covariate partitions creating K random subsets and fitted a linear model for each subset. They constructed a convex function by taking the maximum over these hyperplanes; the hyperplanes induced a new partition, which they used to refit the function. The series of operations repeats until reaching convergence. However, the iterative nature makes the final estimate dependent on the initial

partition. Further, note the estimator is not consistent and there are cases when the algorithm never converges.

Alternatively, [Hannah and Dunson \(2013\)](#) proposed a multi-variate convex adaptive partitioning (CAP) method to estimate locally linear estimates on adaptively selected covariate partitions. CAP uses the upper envelope of the set of local linear estimates to construct a flexible linear hyperplane approximation to the underlying function. The estimator is computationally feasible even for 10,000's of observations and the authors proved its consistency, but its asymptotic rate of convergence is still unknown. [Hannah and Dunson \(2012\)](#) considered a set of ensemble methods such as bagging and smearing which could be applied to the CAP and the Magnani and Boyd (MB) estimators. Bagging subsamples the data set with replacement taking subsample of size n repeated M times [Breiman \(1996\)](#). Each subsample was used to create a new estimate and then the M estimates are averaged. Smearing adds i.i.d. mean zero noise to the observed dependent variable [Breiman \(2000\)](#). A regression model is then fitted to the new noisy data with the observed dependent variables and the M regression estimates are averaged. [Hannah and Dunson \(2012\)](#) implemented smearing and bagging with both CAP and MB. The authors note the significant benefits of bagging and smearing for small sample sizes of 200 observations, and for their large sample sizes of 5,000 observations. In application data, the MB method with either bagging or smearing significantly outperformed (typically by an order of magnitude on a variety of performance criteria) CAP with no augmentations.

[Yagi et al. \(2018\)](#) considers the multi-variate local polynomial kernel estimator with shape constraints. Following their notation, define a set of m points, $\mathbf{x}_1, \dots, \mathbf{x}_m$, for evaluating constraints on the local linear kernel estimator. Recall that $\{(\mathbf{X}_i, Y_i) : i = 1, 2, \dots, n\}$ is the set of observations. [Yagi et al. \(2018\)](#) define the Shape Constrained Kernel-weighted Least Squares (SCKLS) estimator as

$$\begin{aligned} & \underset{\mathbf{a}, \mathbf{b}}{\text{minimize}} && \sum_{i=1}^m \sum_{j=1}^n (Y_j - a_i - (\mathbf{X}_j - \mathbf{x}_i)^\top \mathbf{b}_i)^2 K\left(\frac{\mathbf{X}_j - \mathbf{x}_i}{\mathbf{h}}\right) \\ & \text{subject to} && l(\mathbf{x}_i) \leq \hat{\psi}^{(s)}(\mathbf{x}_i | \mathbf{a}, \mathbf{b}) \leq u(\mathbf{x}_i), \quad i = 1, \dots, m, \end{aligned}$$

where $\mathbf{a} = (a_1, \dots, a_m)^\top$ are functional estimates, $\mathbf{b} = (\mathbf{b}_1^\top, \dots, \mathbf{b}_m^\top)^\top$ are slope estimates at \mathbf{x} , $K(\cdot)$ denotes a product kernel, \mathbf{h} is a vector of bandwidths³, $\hat{\psi}^{(s)}(\mathbf{x}_i | \mathbf{a}, \mathbf{b})$ is the s^{th} derivative of the estimated function $\hat{\psi}$, and $l(\mathbf{x}_i)$, $u(\mathbf{x}_i)$ are the lower and upper bounds, respectively. Specifically, the relationship between $\hat{\psi}$ and the variables \mathbf{a} and \mathbf{b} is $\hat{\psi}(\mathbf{x}_i) = \mathbf{a}_i$ and $\hat{\psi}^{(1)}(\mathbf{x}_i) = \mathbf{b}_i$.

[Yagi et al. \(2018\)](#) showed that SCKLS is consistent and its convergence rate nearly optimal (within a log factor). Unlike other nonparametric estimators, SCKLS uses local information to estimate the functional at any particular point \mathbf{x} , but requires bandwidth selection. Compared to other kernel based shape constrained estimators, such as [Hall and Huang \(2001\)](#) and [Du et al. \(2013\)](#), SCKLS imposes global convexity/concavity by taking the minimum of a set of hyperplanes. A computational complexity analysis implies that SCKLS and [Du et al. \(2013\)](#) method should be similarly difficult to solve because both estimators are solving quadratic objective functions relative to a convex solution spaces. How-

³See [Li and Racine \(2007\)](#) for more details.

ever, practically speaking, SCKLS is usually easier to solve, because the hyperplane structure leads to a sparse constraint matrix, whereas the constraint matrix to restrict the unconstrained kernel estimator is dense making optimization more difficult.

Semiparametric models are common in economics because standard interpretations apply for the linear part. Consider a function f , although unknown, possesses shape properties such as homogeneity, concavity, or monotonicity. [Tripathi \(2000\)](#) considered the standard partially linear model, $Y = \mathbf{X}^\top \boldsymbol{\beta}_0 + f(\mathbf{Z}) + \epsilon$, where Y is the response variable, (\mathbf{X}, \mathbf{Z}) are the covariates, $\boldsymbol{\beta}_0$ is a finite dimensional parameter of interest, and ϵ is an unobserved random Gaussian variable. Tripathi showed that in the class of $n^{1/2}$ consistent regular estimators of the parametric parameters $\boldsymbol{\beta}_0$, the concavity and monotonicity of f do not improve the efficiency of the estimate $\boldsymbol{\beta}_0$ in finite samples. However, homogeneity restrictions on f reduce the lower bounds for the asymptotic variance of the $n^{1/2}$ consistent regular estimators of $\boldsymbol{\beta}_0$.

2.2.3 Application: Production Economics Microeconomic theory, which can be interpreted as shape constraints, provides additional structure for modeling a production or cost function. Consider a production process that uses d different resources to produce a single output, $Y \in \mathbb{R}$. Call the quantity of the resources consumed the inputs and use the notation, $\mathbf{X}_j \in \mathbb{R}^d$. Letting n instances of the production process, each called a production plan, lead to n pairs of input and output data, $\{(\mathbf{X}_j, Y_j)\}_{j=1}^n$. Call the set of all technologically feasible production plans the *production possibilities set* and denote the set as \mathcal{T} . Next, define the production function as

$$Y_j = g_0(\mathbf{X}_j) + \epsilon_j, \quad \text{for } j = 1, \dots, n,$$

where ϵ_j is a random variable satisfying $\mathbf{E}(\epsilon_j | \mathbf{X}_j) = 0$. Here, our primary interest is production function estimation; examples of applications to estimate the dual concept, the cost function, using nonparametric shape constrained estimators include [Beresteanu \(2005\)](#) and [Michaelides et al. \(2015\)](#).

Microeconomic theory often implies basic assumptions, e.g. more input should lead to more output (at least in the input range where the production processes are observed). This particular assumption implies that the production function increases monotonically, specifically

$$\text{if } \mathbf{x}_1 \leq \mathbf{x}_2, \quad \text{then } g_0(\mathbf{x}_1) \leq g_0(\mathbf{x}_2),$$

where the inequality is taken component-wise. Further, for a given output level Y , define the set of input vectors used to produce output level Y as the *input requirement set* (also referred to as the *input set*)

$$V(y) = \{\mathbf{x} : (y, \mathbf{x}) \text{ is in } \mathcal{T}\}.$$

Here, the assumption is an optimal ratio or set of ratios among the inputs exists and any deviation from the optimal ratio requires an increase in other inputs that is more than proportional to the decrease in a particular input. Given two input vectors \mathbf{x}_1 and \mathbf{x}_2 in $V(y)$, then $\lambda \mathbf{x}_1 + (1 - \lambda) \mathbf{x}_2$ is in $V(y)$ for all $0 \leq \lambda \leq 1$. Thus, $V(y)$ is a convex set for any value y ([Varian, 1992](#)). The boundary of the input set is referred to as the input isoquant

$$\text{Isoq } V(y) = \{\mathbf{x} : \mathbf{x} \in V(y), \lambda \mathbf{x} \notin V(y), \lambda < 1\}.$$

The assumption of convex input sets can be strengthened. Define a production function $f(\mathbf{x}) = F(g(\mathbf{x}))$. The production function is *homothetic* if: 1) Scale function $F : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly monotone increasing function, and 2) Core function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is a homogeneous of degree 1 function which implies $g(t\mathbf{x}) = tg(\mathbf{x})$ for all $t > 0$. A production function that does not have this property is referred to as non-homothetic.

The property, *decreasing marginal benefit of inputs*, holds for a variety of production processes. Decreasing marginal benefit of inputs implies that beyond some output level, Y^I , the additional output that can be produced from an additional unit of input decreases as the input level increases:

$$f(\lambda \mathbf{x}) < \lambda f(\mathbf{x}) \text{ for all } \lambda \geq 1 \text{ and } f(\mathbf{x}) \geq y^I.$$

There are two primary reasons for this property. First, for a particular production process certain inputs are well-matched or are the best inputs for that process. *Scarcity of inputs* is the notion that as the scale of production increases, less ideal inputs are used, and so less output per unit of input is achieved. Second, as a production process increases, the related activities are harder to organize or control. Economists call this the *span of control*. A production function with convex input sets that satisfies decreasing marginal benefit of inputs over the entire input space (i.e., $y^I = 0$) is globally concave (Varian, 1984).

For multi-product production, Mundlak (1963) defined a multi-product output vector $\mathbf{y} = (y_1, y_2, \dots, y_q)^\top \in \mathbb{R}_+^q$. For a given input vector \mathbf{x} , define the set of output vectors that can be produced as the *producible output set*:

$$L(\mathbf{x}) = \{\mathbf{y} : (\mathbf{y}, \mathbf{x}) \text{ is in } \mathcal{T}\}.$$

Mundlak (1963) argued that the producible output set $L(\mathbf{x})$ should also be convex. This leads naturally to an implicit *multi-input/multi-output production function*, also called the *transformation function*, defined as:

$$f(\mathbf{y}, \mathbf{x}) = 0.$$

Under the assumption that the input requirement and the producible output sets are convex and that the scaling relationship between inputs and output increases monotonically with the decreasing marginal product, Chambers (1988) shows the multi-input/multi-output production technology \mathcal{T} is globally convex. Further, Chambers (1988) defined the the properties a multi-input/multi-output production technology should satisfy and Kuosmanen and Johnson (2017) provided a nonparametric shape constrained estimator for this technology.

Historically, much of the literature on production functions concerned endogeneity. For example, management would determine input levels with knowledge of its firm-specific characteristics and potentially partial knowledge of random shocks. In the production model, the assumption $\mathbf{E}(\epsilon | \mathbf{X})$ is violated, leading to biased and inconsistent estimates. A variety of solutions have been proposed when the production function is estimated parametrically (Griliches and Mairesse, 1995; Akerberg et al., 2015). Florens et al. (2018) considered estimating a nonparametric shape constrained function using a kernel-based approach

with Landweber–Fridman regularization techniques using instrumental variables to address endogeneity. Research which combines shape constraints with treatment of endogenous variables when estimating production functions is a promising development.

2.3 Single Index Models and Alternative Assumptions to Global Convexity

The desire to maintain convex (linear) input sets while relaxing the concavity assumption for the relationship between the dependent variable and the regressors leads naturally to single index models. Consider the following single index regression model:

$$Y = m_0(\theta_0^\top \mathbf{X}) + \epsilon.$$

Here, $\mathbf{X}_i \in \mathbb{R}^d$ is an observed vector of predictors where $d \geq 1$, $\epsilon_1, \epsilon_2, \dots, \epsilon_n$ satisfy $\mathbf{E}(\epsilon_i | \mathbf{X}_i) = 0$, θ_0 are the parameters of a linear function to project \mathbf{X} to a single dimension, and $m_0 : \mathbb{R}^d \rightarrow \mathbb{R}$ is an unknown link function. The single index model averts the curse of dimensionality typical of nonparametric regression functions with a vector of predictors. This specification states that the link function depends on \mathbf{X} only through a one dimensional projection $\theta_0^\top \mathbf{X}$. Shape restrictions are often placed on the link function m_0 . [Balabdaoui et al. \(2016\)](#), who studied this model under a monotonicity constraint on m_0 , proved consistency of the least squares estimator and [Balabdaoui et al. \(2017\)](#) proved that the least squares estimator achieves the $n^{-1/2}$ rate of convergence. [Kuchibhotla et al. \(2017\)](#) considered both a Lipschitz constrained least squares estimator and the penalized least squares estimator and found similar results for consistency and rate of convergence. The single index model adds structure to the estimator, which can be useful particularly when the data are limited. However, assuming linear substitution between inputs or goods, $\theta_0^\top \mathbf{X}$, can be overly restrictive for many production or utility models.

2.3.1 Application: Production Functions Continued Convex input sets are typically a maintained assumption in production economics, but alternative assumptions are available for the scaling law (or the relationship between output and expanding input levels). While decreasing marginal product is a common characteristic for large firms, economists often assume an increase in marginal product for production at a small scale, i.e. increasing returns to scale. Thus, a proportional increase in inputs leads to a more than proportional increase in output

$$f(\lambda \mathbf{x}) > \lambda f(\mathbf{x}) \text{ for all } \lambda \geq 1.$$

The primary reasons for the increasing returns to scale are the benefit of specialization and the reduction in change-over time when switching between components of the production process.

[Frisch \(1964\)](#) proposed the Regular Ultra Passum (RUP) production law. The law outlines when a firm is operating at a small scale size it can achieve significant increases in output for incremental increase in input through specialization and learning. In contrast, as the scale size becomes larger, a firm tends to face scarcity of ideal production inputs and challenges related to increasing span of control, thus the marginal benefits of additional inputs decrease. Based on these concepts,

define the elasticity of scale,⁴ $\epsilon(\mathbf{x})$, relative to a production function $f(\mathbf{x})$:

$$(2.5) \quad \epsilon(\mathbf{x}) = \sum_{k=1}^d \frac{\partial f(\mathbf{x})}{\partial x_k} \frac{x_k}{f(\mathbf{x})},$$

and define the RUP law as follows.

DEFINITION 2 (Førsund and Hjalmarsson (2004)). *A production function $f(\mathbf{x})$ obeys the Regular Ultra Passum Law if $\partial\epsilon(\mathbf{x})/\partial x_k < 0$, and there exist input vectors \mathbf{x}_a and \mathbf{x}_b where $\mathbf{x}_b \geq \mathbf{x}_a$ component-wise such that $\epsilon(\mathbf{x}_a) > 1$ while $\epsilon(\mathbf{x}_b) < 1$.*⁵

While the RUP law is often not referred to by name, most introductory microeconomic textbooks introduce the concept (Perloff, 2018). Further, researchers commonly assume that along a given ray, there exists only a single inflection point; however, neither Førsund and Hjalmarsson (2004) nor Frisch (1964) definition rules out this possibility. To make this concept rigorous, Yagi et al. (2018) defined an S-shape function.

DEFINITION 3 (Yagi et al. (2018)). *For any vector $\mathbf{v} \in \mathbb{R}_+^d$ and the associated ray from the origin in input space $\alpha\mathbf{v}$ with $\alpha > 0$, a production function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is S-shaped if there exists an x^* such that $\nabla_{\mathbf{v}}^2 f(\alpha\mathbf{v}) > 0$ for $\alpha\mathbf{v} < \mathbf{x}^*$, and $\nabla_{\mathbf{v}}^2 f(\alpha\mathbf{v}) < 0$ for $\alpha\mathbf{v} > \mathbf{x}^*$, where $\nabla_{\mathbf{v}}^2 f$ is the directional second derivative of f along \mathbf{v} . This implies that for any ray from the origin in the direction \mathbf{v} , there exists a single inflection point \mathbf{x}^* that $\nabla_{\mathbf{v}}^2 f(\mathbf{x}^*) = 0$.*

In defining an S-shape production function, Yagi et al. (2018) derived the relationship to the long standing concept of the RUP law. Specifically, if a production function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is twice-differentiable, monotonically increasing, satisfies the RUP law, and has a single inflection point, then f is S-shaped.

Yagi et al. (2018) proposed an estimation algorithm for a non-homothetic production function satisfying both the S-shape definition and input convexity without any further structural assumptions. The algorithm has two steps: (1) Estimate input isoquants for a set of output levels, and (2) estimate S-shape functions on a set of rays from the origin. A CNLS-based estimator is used for isoquant estimation and a SCKLS-based estimators for the S-shape estimation. While this estimator has a similar flavor as the single index model, relaxing the parameter structure for aggregating the vector of inputs necessitates a two step procedure. In the non-homothetic case, the performance of the estimator can be improved by iterating between the two-steps. However, if the production function is homothetic, the input isoquant can be estimated for just a single output level as in Hwangbo et al. (2015).

⁴This variable, formerly referred to as the passum coefficient in the seminal work of Frisch (1964), is now commonly referred to as the elasticity of scale.

⁵Note that this definition of the RUP law is slightly adapted from Førsund and Hjalmarsson (2004) for clarity. Førsund and Hjalmarsson (2004) definition generalizes Frisch (1964) original definition by not requiring the passum coefficient to go below 0 implying congestion. This generalization also allows for a monotonically increasing production function. Further, a concave production function nests within this definition.

In the future, the availability of shape constrained nonparametric estimation techniques could allow economists to develop alternative theories of production and validate them empirically.

3. SEQUENTIAL DECISION MAKING

Making sequential decisions under uncertainty, which has been studied extensively in both economics and operations research, is formalized under *dynamic programming*, *optimal control*, *Markov decision processes*, and other topics (Stokey et al., 1989; Kamien and Schwartz, 1981; Puterman, 1994). The typical setting consists of a decision-maker who alternates between making decisions and observing new information, to inform future decisions:

choose decision \rightsquigarrow observe information \rightsquigarrow choose next decision $\rightsquigarrow \dots$

In a sequential setting where stochastic information is revealed over time, the decision maker needs to find the *optimal policy* that prescribes a decision for every possible “state-of-the-world” i.e., every outcome of the stochastic information process. The following example, illustrates the well-known problem of multi-stage inventory management (Clark and Scarf, 1960; Scarf, 1960; Porteus, 2002).

Consider a firm managing its inventory control policy over a finite horizon of T periods. At time period t , the decision maker observes the current inventory state s_t and places an order for x_t additional units. Between time t and time $t+1$, a random demand D_{t+1} , independent of the past, is realized. The cost function for period t is given by

$$(3.1) \quad c_t(s_t, x_t) = cx_t + \mathbf{E}[h(s_t + x_t - D_{t+1})^+ + b(D_{t+1} - s_t - x_t)^+],$$

where c is the ordering cost, h is the holding cost, and b is the backlogging cost (i.e., cost per unit of unsatisfied demand). The inventory position at $t+1$ is given by $s_{t+1} = s_t + x_t - D_{t+1}$, where $s_t < 0$ represents unsatisfied or backlogged demand. The decision maker needs to determine the optimal inventory ordering policy π_t^* (a function mapping inventory states s_t to order quantities x_t) that minimizes the expected cumulative cost $\mathbf{E}[\sum_{t=0}^{T-1} c_t(s_t, \pi_t(s_t))]$.

At every period t , the decision maker needs to determine an order quantity x_t for each possible inventory state s_t (a scalar quantity). This model is known as a *finite-horizon Markov decision process* (MDP) (Puterman, 1994). If s_t takes on a finite number of values and the expected value is easy to compute (e.g., in the case where order quantities are integer-valued and demands are integer-valued and bounded), then the optimal decision in stage t , state s , denoted $\pi_t^*(s)$, can be computed by simply enumerating the inventory states and then applying a standard dynamic programming method (Bertsekas, 2012; Puterman, 1994). This amounts to solving a series of recursive equations to compute the so-called *optimal value functions*. Let $V_T(s) \equiv 0$ and for each $t < T$, define the optimal value function

$$(3.2) \quad V_t(s_t) = \min_{x_t} \{c_t(s_t, x_t) + \mathbf{E}[V_{t+1}(s_t + x_t - D_{t+1})]\}.$$

Here, the optimal decision in state s_t at time t is an x_t that achieves the minimum of the right-hand-side of (3.2). Thus, solving for the optimal policy π_t^* for each t

under the dynamic programming framework depends on the decision maker’s ability to compute the value functions, V_t . Although the recursive equations may seem simple, realistic instances of the inventory control problem can quickly become intractable for the enumeration-based dynamic programming method discussed above. For example,

- Consider the case of a *multi-product* inventory system, as described in [Evans \(1967\)](#) or [Aviv and Federgruen \(2001\)](#). In this setting, the decision maker needs to track the inventory states for all products; thus, s_t becomes multi-dimensional. Even if the number of inventory states per product is enumerable, the number of states across all products grows exponentially with the number of products; this is an example of the “curse of dimensionality.”
- The assumption that D_{t+1} is independent of the past is called *stage-wise independence* ([Pereira and Pinto, 1991](#); [Shapiro, 2011](#)). In practical applications, the distribution of the demand could depend on factors such as weather, previous demands, or market conditions represented by i_t . Consequently, the optimal policy depends on both s_t and i_t , and is written as $\pi_t^*(s_t, i_t)$. Computational difficulties easily arise when i_t is multi-dimensional, again due to the curse of dimensionality.
- An unknown distribution for D_{t+1} implies the inability to compute the expectation in (3.1). This limitation, which prevents the application of standard dynamic programming techniques, forces the decision maker to use sample-based methods, usually historical data or a simulator (generative model).

Approximate dynamic programming (ADP) and reinforcement learning (RL) refer to a set of methodologies and algorithms for approximately solving complex sequential decision problems when the state space is large and/or parts of the system are unknown ([Bertsekas and Tsitsiklis, 1996](#); [Sutton and Barto, 1998](#); [Bertsekas, 2012](#); [Powell, 2011](#)). For general large-scale problems, *unstructured approximations* including value iteration with linear approximations (i.e., using basis functions) ([De Farias and Van Roy, 2000](#); [Tsitsiklis and Roy, 1996](#); [Tsitsiklis and Van Roy, 1999](#); [Geramifard et al., 2013](#)), approximate linear programming ([De Farias and Van Roy, 2003](#); [De Farias and Van Roy, 2004](#); [Desai et al., 2012a](#)), and nonparametric methods are used ([Ormoneit and Sen, 2002](#); [Bhat et al., 2012](#)). Recently, RL with deep neural networks has become popular ([Mnih et al., 2015](#); [Silver et al., 2016](#)). *Structured approximations* can be incorporated when properties can be identified *a priori*.

The next section reviews the uses of shape constraints to enforce structure in the approximations used in ADP and RL. The primary focus is on convexity of the value function, a particularly well-studied shape constraint in sequential decision problems ([Pereira and Pinto, 1991](#); [Godfrey and Powell, 2001](#); [Philpott and Guan, 2008](#); [Nascimento and Powell, 2009, 2010](#)), and the secondary focus is monotonicity ([Papadaki and Powell, 2002](#); [Jiang and Powell, 2015](#)), which is useful when convexity is not available. We conclude with a brief review of how *policy structure* can be exploited ([Kunnumkal and Topaloglu, 2008b](#); [Huh and Rusmevichientong, 2009](#); [Zhang et al., 2017](#)). The notation we use will largely follow the standards of the literature.

3.1 Convexity of the Value Function

In many problems, the value function V_t is *convex* in the state variable, or, at least in certain dimensions of the state variable, for each t . One benefit of the property of convexity is the potential for algorithms to exploit this structure. Below, we survey several widely used and powerful methodologies.

3.1.1 Stochastic Decomposition We begin with the *two-stage linear program with recourse* (Birge and Louveaux, 2011). The *recourse* decisions are made in the second (and final stage) after uncertainties have been realized. The interpretation is that the “mistakes” made by the first-stage decision (e.g., inventory shortage) can be corrected in the second stage. Let \mathcal{X} be a convex, polyhedral set. Using the notation of Hige and Sen (1991), formulate the problem as

$$\begin{aligned} & \underset{x}{\text{minimize}} && f(x) = c^\top x + \mathbf{E}\{h(x, \tilde{\omega})\} \\ & \text{subject to} && x \in \mathcal{X}, \end{aligned}$$

where $h(x, \omega)$ is the optimal objective function value of the second-stage problem

$$\begin{aligned} & \underset{y}{\text{minimize}} && g^\top y \\ & \text{subject to} && Wy = \omega - Tx \\ & && y \geq 0. \end{aligned}$$

Here, c is a cost vector and x is the first-stage decision, which must be chosen before the realization of a random variable $\tilde{\omega}$. The second-stage *recourse costs* are given by g and the associated decision is y . The *recourse matrix* is W and $\omega - Tx$ is discrepancy that is accounted for in the second stage by y . Assume that value function h is finite. To illustrate the formulation, in a production problem, x may represent raw material order quantities from a supplier (at cost c) to produce Tx units of a product; ω is the realized demand and $\omega - Tx$ is a shortage; and y are the raw material order quantities from an emergency supplier (at costs g), which can produce Wy units of the product.

The stochastic decomposition (SD) algorithm proposed by Hige and Sen (1991) exploits the *convexity and piecewise-linearity* of $h(x, \omega)$ by combining ideas from Bender’s decomposition and stochastic approximation (Kushner and Yin, 2003). The algorithm constructs a necessarily convex approximation of the objective function f by defining it as the maximum of iteratively computed hyperplanes. The steps of the stochastic decomposition method are summarized below.

- (1) *Subproblem.* Given x^k on iteration k , compute the solution to the second-stage problem for a single sample ω^k generated from the distribution of $\tilde{\omega}$ by using a dual formulation, with dual variables π

$$h(x^k, \omega^k) = \max_{\pi} \{ \pi^\top (\omega^k - Tx^k) \mid W^\top \pi \leq g \}.$$

- (2) *Compute Cut.* Using all subproblem solutions generated until iteration k , estimate the support of $f(x)$ via an affine function $\alpha_t^k + (\beta_t^k + c)x$. This affine function is called a *cut*. See (Hige and Sen, 1991, Section 2.2) for the details.

- (3) *Update Old Cuts.* Reduce the influence of the old cuts generated in iterations $t < k$ via the updates

$$\alpha_t^k = (k-1)\alpha_t^{k-1}/k \quad \text{and} \quad \beta_t^k = (k-1)\beta_t^{k-1}/k.$$

- (4) *Update Convex Approximation.* Compute the piecewise-linear and convex approximation of the objective function $f(x)$ as a maximum of affine functions:

$$f_k(x) = \max_{t \leq k} \{ \alpha_t^k + (\beta_t^k + c)x \}$$

The iterate x^{k+1} is a solution to an *approximate* first-stage problem given by $\max_{x \in \mathcal{X}} f_k(x)$.

THEOREM 3.1 (Higle and Sen (1991)). *There exists a subsequence of $\{x^k\}$ generated by SD, such that every accumulation point of the subsequence is an optimal solution, with probability one.*

The formal statement of the SD algorithm also uses the concept of an incumbent solution, which roughly speaking, is an iterate which achieves a low objective value and is revisited by the algorithm. Such an implementation improves the method's empirical performance and does not affect the convergence guarantee. However, despite asymptotic guarantees, stopping rules are critical for practical implementations (see, e.g., Higle and Sen (1996); Mak et al. (1999)). The original SD algorithm was designed only for two-stage problems, but more recently, Sen and Zhou (2014) proposed a regularized, multi-stage extension.

3.1.2 Stochastic Dual Dynamic Programming The *stochastic dual dynamic programming* (SDDP) algorithm was first proposed by Pereira and Pinto (1991). SDDP was proposed before SD, but both are based on the idea of iteratively generating cuts to approximate a piecewise-linear convex value function. The difference is that SDDP does not use stochastic approximation (i.e., the phase-out step), but requires the subproblems to be solved for every scenario in each iteration. SDDP is more general in the sense that it was directly proposed for multistage problems.

We introduce the multistage stochastic linear programming model using the notation of Philpott and Guan (2008), while noting that it is a direct extension of the two-stage model discussed in the previous section. Let Ω_t be a finite set of random outcomes for state t , where outcome ω_{ti} has probability p_{ti} . The random variables $\omega_t \in \Omega_t$ are independent across time (termed *stagewise independence*). Let $\mathcal{Q}_{T+1} \equiv 0$ and define $\mathcal{Q}_t(x_{t-1}) = \sum_i p_{ti} \mathcal{Q}_t(x_{t-1}, \omega_{ti})$, where

$$(3.3) \quad \begin{aligned} \mathcal{Q}_t(x_{t-1}, \omega_{ti}) &= \min_{x_t} c_t x_t + \mathcal{Q}_{t+1}(x_t) \\ &\text{subject to } A_t x_t = \omega_t - B_{t-1} x_{t-1} \\ &\quad x_t \geq 0, \end{aligned}$$

and write the first-stage problem as

$$\begin{aligned} Q_1 &= \min_{x_1} c_1 x_1 + Q_2(x_1) \\ &\text{subject to } A_1 x_1 = b_1 \\ &\quad x_1 \geq 0. \end{aligned}$$

Once again, the value functions $Q_t(x)$ are piecewise-linear convex (Pereira and Pinto, 1991), so the approximation used in SDDP is a maximum of hyperplanes. Let Q_t^k be the approximation of Q_t at iteration k and suppose $Q_{T+1}^k = Q_{T+1}$ for all k . The main steps of iteration k of the algorithm are as follows.

- (1) *Forward Pass.* For each time period $t = 1, 2, \dots, T$, select a *trial decision* \hat{x}_t^k . There are a variety of ways to do this, but the original proposal from Pereira and Pinto (1991) is to simulate one sample of the decisions generated by the current policy, i.e., the one induced by solving (3.3) where Q_{t+1} is replaced with the current approximate value function Q_{t+1}^{k-1} .
- (2) *Backward Pass.* SDDP now moves backward, starting from $t = T$. At the trial decision \hat{x}_t^k , loop through all possible outcomes of the random variable ω_t and for each outcome ω_{ti} , solve the problem (3.3) with Q_{t+1}^k replacing Q_{t+1} and ω_{ti} replacing ω_t . Compute a cut and use it to update the approximation for time t , resulting in Q_t^k (which is used in the subproblem solved in at $t - 1$). Note that the convexity of Q_t is again exploited by approximating it as the maximum of a series of affine cuts.

The convergence of the SDDP algorithm has been studied by Linowsky and Philpott (2005) and further generalized by Philpott and Guan (2008). Note that the improvement of the convergence result here compared to the one for SD is because all scenarios are being considered when computing the cuts.

THEOREM 3.2 (Philpott and Guan (2008)). *Under some technical assumptions on the forward pass, the SDDP algorithm converges to an optimal solution of the first-stage problem in a finite number of iterations.*

Two algorithms closely related to SDDP are the *cutting plane and partial sampling* algorithm of Chen and Powell (1999) and the *abridged nested decomposition* algorithm of Donohue and Birge (2006). Notably, the convergence of both methods follows from Philpott and Guan (2008). The stagewise independence assumption of the random process was relaxed in Löhndorf et al. (2013) and Asamov and Powell (2018). Asamov and Powell (2018) also proposed a regularized version of SDDP to improve performance. Extensions of SDDP for risk-averse problems were explored in Philpott and de Matos (2012), Shapiro et al. (2013), and Philpott et al. (2013). All of these methods enforce a convex “shape constraint” on the value function approximation via a piecewise-linear function.

EXAMPLE (Hydrothermal Planning). *Hydrothermal operations planning has a long history in the operations research literature (Pereira and Pinto, 1991; Pritchard et al., 2005; Philpott and de Matos, 2012; Shapiro et al., 2013; Maceira et al., 2015). The objective is to find an operational strategy that satisfies energy demand at every location in the system while achieving minimal expected cost. Below, we summarize the major features.*

- *The overall system contains a set of reservoirs whose storage levels are tracked by the state variable. From period to period, the reservoir is subjected to stochastic inflows and potential losses due to evaporation.*
- *Water stored in the reservoirs can be used for (free) energy production. Thermal generators used to complement the hydroelectric production are*

expensive to operate, i.e., each generator has an associated generation cost function.

- Energy can be interchanged between two locations via transmission lines. This feature emphasizes the influence of the underlying network structure.

Discretization and standard dynamic programming cannot be used, because of the increased number of states as the number of reservoirs increases. Instead, SDDP is used to provide approximations of the optimal policy for problems ranging in size from 22 reservoirs (Pereira and Pinto, 1991) to 69 reservoirs (Shapiro et al., 2013).

3.1.3 Separable, Projective Approximation Routine As proposed by Powell et al. (2004), the *separable, projective approximation routine* (SPAR), was motivated by the need for integer solutions in certain applications (e.g., consider the case of ordering a discrete number of items per period). The term “separable approximation” means that a multi-dimensional convex function is approximated by a sum of one-dimensional convex functions, which unfortunately implies that the algorithm does not converge to an optimal solution in general. However, its empirical performance is strong. Nascimento and Powell (2013) extended theory in Powell et al. (2004) to a multistage setting for the (one-dimensional convex) problem of lagged asset acquisition. The SPAR algorithm relies on stochastic approximation theory and bears a strong resemblance to the Q -learning algorithm (Watkins and Dayan, 1992) in the reinforcement learning literature (Sutton and Barto, 1998). Below, we describe a single period setting.

Using the notation of Powell et al. (2004), consider maximization of a piecewise-linear concave function $\bar{f}(x) = \mathbf{E}\{f(x, \omega)\}$ with integer breakpoints on a domain $[0, M]$. If $l \leq x < l + 1$, write function \bar{f} as

$$\bar{f}(x) = \bar{f}(0) + \sum_{s=1}^l \bar{v}_s + \bar{v}_{l+1} \cdot (x - l),$$

where $\bar{v}_s = \bar{f}(s) - \bar{f}(s - 1)$ are the slopes of \bar{f} . The idea of the SPAR algorithm is to iteratively construct approximations of \bar{f} using noisy observations while employing a concavity preservation step to ensure that \bar{v}_s are nonincreasing in s . The steps of iteration k , where the current approximation of the true slope \bar{v}_s^k is v_s^k are as follows.

- (1) *Noisy Observation.* Sample a “state” s^k from $\{1, 2, \dots, M\}$ and observe (from data or a simulator) an unbiased estimate η^k of the slope \bar{v}_{s^k} .
- (2) *Update Approximation.* Compute an intermediate approximation z^k , with $z_{s^k}^k = (1 - \alpha_k)v_{s^k}^k + \alpha_k \eta^k$ and $z_s^k = v_s^k$ for all other s . This step simply smoothes the new observation with the current estimate.
- (3) *Enforce Concavity.* Let V be the set of vectors \bar{v} that represent concave functions \bar{f} , i.e., those that satisfy $\bar{v}_{s+1} \leq \bar{v}_s$ for each s . The SPAR algorithm enforces concavity by projecting z^k to the set V

$$v^{k+1} = \operatorname{argmin}\{\|v - z^k\|^2 : v \in V\},$$

which can be computed via a straightforward procedure. This step can be thought of as correcting any nonconcavity introduced by the noisy update.

THEOREM 3.3 (Powell et al. (2004)). *Under some technical conditions, the approximations v^k produced by the SPAR method converge almost surely to \bar{v} .*

The SPAR procedure can be used within an approximate dynamic programming setting for multistage decision making, as analyzed in Nascimento and Powell (2009). Related algorithms that make use of stochastic approximation to update piecewise linear approximations can be found in Cheung and Powell (2000) and Godfrey and Powell (2001). This class of techniques is more amenable to the distribution-free setting, where sampling can be done with respect to the true distribution or from real data, whereas SDDP requires either a pre-sampled model (with a finite number of scenarios per stage) or the assumption that the true distribution has finite support and is known (Shapiro, 2011). SPAR's limitations are its requirement for a separable approximation for multi-dimensional problems, and the unavailability of lower bounds for the minimization case. See Asamov et al. (2016) for a thorough empirical comparison of SPAR versus SDDP for the case of optimizing grid-level energy storage; the authors concluded that each algorithm has benefits in certain contexts.

EXAMPLE (Cash Balance). *The mutual fund cash balance problem (Nascimento and Powell, 2010), is an updated version of the stochastic cash balance problem (Neave, 1970). Consider the decision problem of a mutual fund manager, who requires a policy to decide the cash quantity of the fund's assets kept in each period. Several trade-offs need to be considered.*

- *A mutual fund needs to meet shareholder redemptions. There is redemption demand from retail (small) investors and institutional (large) investors, who are treated slightly differently. If there is not enough cash on hand to satisfy the redemption of a large investor, the fund manager takes out short-term loans to immediately meet the demand.*
- *In the case of a retail investor, the fund manager pays a cost to liquidate a portion of the individual's illiquid assets. Generally, the cost is smaller than the interest on the short-term loans.*
- *If too much cash is kept on hand, the fund manager foregoes the portfolio return on the portion of excess cash.*

The authors formulated the above problem as an MDP and proved that the optimal value function was piecewise-linear concave in the cash level for each environmental state. Given that the problem is one-dimensional and that the two exogenous stochastic processes (portfolio return and interest rate) have unknown distributions, SPAR is an ideal solution algorithm.

3.2 Monotonicity

Some studies have exploited *monotone structure* in sequential decision making. In the scalar case, the discussion above shows that enforcing concavity is equivalent to enforcing monotonicity in the slopes v_s^k . Similar methods have been designed for this setting. See Papadaki and Powell (2002) for an example of the SPAR-like procedure adapted to exploit monotonicity in the stochastic batch service problem. Kunnumkal and Topaloglu (2008a) proposed a *monotone Q-learning* for the setting of scalar state variables in an *infinite horizon setting*; for convergence analysis, it requires a different proof technique from the finite horizon

setting of [Nascimento and Powell \(2013\)](#). However, for a multi-dimensional state, [Kunnumkal and Topaloglu \(2008a\)](#) suggest arbitrarily selecting a dimension in which monotonicity is enforced.

[Jiang and Powell \(2015\)](#) proposed *monotone-ADP*, extending previous work to the multi-dimensional state space setting. More generally, the algorithm applies when monotonicity is preserved over *partially ordered* states. Consider a dynamic programming setting similar to (3.2), where the value functions $V_t(s)$ satisfy

$$V_t(s) \leq V_t(s') \quad \text{for all } s \preceq s',$$

where \preceq is a partial order over the state space. Monotone-ADP is useful in the setting of $s = (x_1, x_2, \dots, x_d)$ and $s' = (x'_1, x'_2, \dots, x'_d)$ with monotonicity in the value function whenever $x_i \leq x'_i$ for all i . The steps of the algorithm are analogous to those of SPAR, except in a multistage setting. For example, in iteration k , there is a state s_t^k , an updated value z_t^k , and a current value function approximation V_t^k . Here, we focus on a new operator Π_M , which depends on s_t^k , z_t^k , and V_t^k , where there is an arbitrary state s and the output is the updated value of state s after accounting for the new observation at s_t^k and the monotone structure. Define it as

$$\Pi_M(s_t^k, z_t^k, V_t^k)(s) = \begin{cases} z_t^k & \text{if } s = s_t^k, \\ z_t^k \vee V_t^k(s) & \text{if } s_t^k \preceq s, s \neq s_t^k, \\ z_t^k \wedge V_t^k(s) & \text{if } s_t^k \succeq s, s \neq s_t^k, \\ V_t(s) & \text{otherwise,} \end{cases}$$

where $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$. The first condition says that the value of s_t^k is updated to z_t^k regardless of the other values, which is a departure from SPAR. The second condition says that if s is greater than s_t^k , then its value should be the maximum of z_t^k and V_t^k under the monotonicity condition. The third condition similarly covers the case of s less than s_t^k . The fourth condition leaves the ‘‘incomparable’’ states unchanged, because there is only a partial order over the states.

THEOREM 3.4 ([Jiang and Powell \(2015\)](#)). *Under some technical assumptions, the approximation $V_t^k(s)$ generated by monotone-ADP converges to the value function $V_t(s)$ almost surely for each stage t and state s .*

Since the theoretical basis of the monotone-ADP algorithm is stochastic approximation, this theorem emphasizes that the monotone shape constraint does not affect its convergence properties.

EXAMPLE (Optimal Stopping). *Perhaps the most fundamental problem class in sequential decision making concerns the question of optimal stopping or optimal replacement ([Pierskalla and Voelker, 1976](#); [Rust, 1987](#); [Tsitsiklis and Van Roy, 1999](#); [Kurt and Maillart, 2009](#); [Kurt and Kharoufeh, 2010](#); [Desai et al., 2012b](#)). The trade-off in the case of optimal stopping is whether the decision maker should accept the reward now (e.g., sell a house), or wait until a future period (e.g., wait for a higher offer). Similarly, optimal replacement problems model the trade-off between the cost of replacement now versus the possibility of failure in the future.*

Although such problems do not feature convex value functions, they often contain many exogenous information states, which impose significant computational challenges. Fortunately, monotonicity can sometimes hold. For example, Kurt and Maillart (2009) and Kurt and Kharoufeh (2010) showed that under certain conditions, the value function of an optimal replacement model was nondecreasing in both the system state (health of the system to be replaced) and the environmental state, i.e. an opportunity for monotone shape constraints to be utilized was provided. Jiang and Powell (2015), who applied the monotone-ADP algorithm to problems of this type, showed significant empirical improvements when monotonicity was enforced.

3.3 Policy Structure

The final goal of sequential decision making problems is to discover the *optimal policy*. Computing the optimal value function is one approach, but it is also possible to directly search for the optimal policy $\pi_t^*(s)$. Consider the well-known policy gradient methods of reinforcement learning (Sutton et al., 2000). In operations research, the structure of optimal policies are often known *a priori* for many foundational models. Below, we give an example of using shape constraints in the policy space.

It is well-known that the optimal policy to the inventory control application discussed above is of *basestock form* (for example see Porteus (2002)), meaning that optimal basestock levels $\{r_t\}$ exist such that if the current inventory level is s_t , then it is optimal to place an order for $x_t = [r_t - s_t]^+$ units of inventory. In other words, if s_t is below r_t , the decision maker orders up to r_t , and if s_t is above r_t , it does not order.

Kunnumkal and Topaloglu (2008b) proposed an algorithm that directly searches within the class of all basestock policies, in effect imposing a “basestock shape constraint” throughout the search. The authors considered a stochastic approximation approach, where $\{r_t^k\}$ denotes the estimated basestock levels at iteration k . The two steps are as follows.

- (1) *Demand Observations*. Observe either from data or within a simulated setting, a trajectory of demand realizations

$$\mathbf{D}^k = (D_1^k, D_2^k, \dots, D_T^k).$$

Note that demand is exogenous to the inventory system.

- (2) *Basestock Update*. Using the current basestock levels $\{r_t^k\}$ and the demand observations \mathbf{D}^k , compute an estimated basestock adjustment Δ_t^k to improve the performance of the basestock policy. The update is given by

$$r_t^{k+1} = r_t^k - \alpha^k \Delta_t^k$$

for some *stepsize* or *learning-rate* α^k .

Because no value function approximation is stored, the calculation of Δ_t^k is non-trivial and requires a novel recursive computation derived from the Bellman equation (Kunnumkal and Topaloglu, 2008b, Section 3).

THEOREM 3.5 (Kunnumkal and Topaloglu (2008b)). *Under various technical conditions, the sequence of basestock policies generated by the method above is asymptotically optimal: $r_t^k \rightarrow r_t$ almost surely for each t .*

A basestock shape constraint was also used in the online convex optimization approach of [Huh and Rusmevichientong \(2009\)](#). A related paper by [Zhang et al. \(2017\)](#), used the basestock shape constraint in a perishable inventory setting to discover good, but not necessarily optimal, policies.

4. DISCUSSION

The economics and operations research communities continue to propose, test, and debate new applications of shape constrained estimation. This paper briefly surveys several significant applications of shape constraints, including revealed preferences, production economics, and several operational problems that involve sequential decision making. To our knowledge, this survey is the first to review the operations research literature on shape constraints.

Shape constraints, including monotonicity, convexity/concavity, and S-shapes, have attracted research from both practical and theoretical perspectives. The additional structure provided by imposing shape constraints allows for estimation with smaller samples, and in the case of sequential decision making applications, structured versions of approximate dynamic programming algorithms enables practitioners to address large-scale problems. We expect that new methodologies will be developed to address other important problem structures, which will in turn facilitate the growth of the fields of economics and operations research.

ACKNOWLEDGEMENTS

We are grateful to Bodhisattva Sen and Richard Samworth for inviting us to write this paper. Andrew thanks Osaka University for its financial support. Any mistakes are our own. The views expressed herein are those of the authors.

REFERENCES

- Akerberg, D. A., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica* 83(6), 2411–2451.
- Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International economic review* 8(1), 67–77.
- Afriat, S. N. (1972). Efficiency estimation of production functions. *International Economic Review* 13(3), 568–598.
- Ait-Sahalia, Y. and J. Duarte (2003). Nonparametric option pricing under shape restrictions. *Journal of Econometrics* 116(1), 9–47.
- Alizadeh, F. (2006). Semidefinite and second-order cone programming and their application to shape-constrained regression and density estimation. In *Models, Methods, and Applications for Innovative Decision Making*, pp. 37–65.
- Allon, G., M. Beenstock, S. Hackman, U. Passy, and A. Shapiro (2007). Nonparametric estimation of concave production technologies by entropic methods. *Journal of Applied Econometrics* 22(4), 795–816.
- Asamov, T. and W. B. Powell (2018). Regularized decomposition of high-dimensional multistage stochastic programs with Markov uncertainty. *SIAM Journal on Optimization* 28(1), 575–595.
- Asamov, T., D. F. Salas, and W. B. Powell (2016). SDDP vs. ADP: The effect of dimensionality in multistage stochastic optimization for grid level energy storage. arXiv preprint arXiv:1605.01521.
- Aviv, Y. and A. Federgruen (2001). Capacitated multi-item inventory systems with random and seasonally fluctuating demands: implications for postponement strategies. *Management science* 47(4), 512–531.
- Balabdaoui, F., C. Durot, and H. Jankowski (2016). Least squares estimation in the monotone single index model. arXiv preprint arXiv:1610.06026.
- Balabdaoui, F., P. Groeneboom, and K. Hendrickx (2017). Score estimation in the monotone single index model. arXiv preprint arXiv:1712.05593.
- Banker, R. D. and A. Maindiratta (1992). Maximum likelihood estimation of monotone and concave production frontiers. *Journal of Productivity Analysis* 3(4), 401–415.
- Barlow, R. E., D. J. Bartholomew, J. M. Bremner, and H. D. Brunk (1972). *Statistical Inference Under Order Restrictions; The Theory and Application of Isotonic Regression*. Wiley, New York.
- Beresteanu, A. (2005). Nonparametric analysis of cost complementarities in the telecommunications industry. *RAND Journal of Economics* 36(4), 870–889.
- Beresteanu, A. et al. (2007). Nonparametric estimation of regression functions under restrictions on partial derivatives. Preprint. <http://www.econ.duke.edu/arie/shape.pdf>.
- Bertsekas, D. P. (1999). *Nonlinear programming*. Athena scientific Belmont.
- Bertsekas, D. P. (2012). *Dynamic programming and optimal control, Vol. II: Approximate dynamic programming* (4th ed.). Belmont, MA: Athena Scientific.
- Bertsekas, D. P. and J. N. Tsitsiklis (1996). *Neuro-dynamic programming*. Belmont, MA: Athena Scientific.
- Bhat, N., V. Farias, and C. C. Moallemi (2012). Non-parametric approximate dynamic programming via the kernel method. In *Advances in Neural Information Processing Systems 25*, pp. 386–394.
- Bhattacharya, D. (2016). Applied welfare analysis for discrete choice with interval-data on income. Working Paper.
- Birge, J. R. and F. Louveaux (2011). *Introduction to stochastic programming* (2nd ed.). Springer Science & Business Media.
- Blundell, R., M. Browning, and I. Crawford (2008). Best nonparametric bounds on demand responses. *Econometrica* 76(6), 1227–1262.
- Blundell, R. W., M. Browning, and I. A. Crawford (2003). Nonparametric engel curves and revealed preference. *Econometrica* 71(1), 205–240.
- Breiman, L. (1996). Bagging predictors. *Machine learning* 24(2), 123–140.
- Breiman, L. (2000). Randomizing outputs to increase prediction accuracy. *Machine Learning* 40(3), 229–242.
- Brown, D. J. and R. L. Matzkin (1996). Testable restrictions on the equilibrium manifold. *Econometrica: Journal of the Econometric Society* 64(6), 1249–1262.
- Brunk, H. (1970). Estimation of isotonic regression. In *Nonparametric Techniques in Statistical*

- Inference*, pp. 177–195. Cambridge University Press.
- Chambers, R. G. (1988). *Applied production analysis: a dual approach*. Cambridge University Press.
- Chatterjee, S., A. Guntuboyina, B. Sen, et al. (2015). On risk bounds in isotonic and other shape restricted regression problems. *The Annals of Statistics* 43(4), 1774–1800.
- Chen, Y. and R. J. Samworth (2016). Generalized additive and index models with shape constraints. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78(4), 729–754.
- Chen, Z.-L. and W. B. Powell (1999). Convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse. *Journal of Optimization Theory and Applications* 102(3), 497–524.
- Cherchye, L., B. De Rock, and F. Vermeulen (2007). The collective model of household consumption: a nonparametric characterization. *Econometrica* 75(2), 553–574.
- Chernozhukov, V., W. K. Newey, and A. Santos (2015). Constrained conditional moment restriction models. arXiv preprint [arXiv:1509.06311](https://arxiv.org/abs/1509.06311).
- Chetverikov, D., A. Santos, and A. M. Shaikh (2018). The econometrics of shape restrictions. *Annual Review of Economics* 10(In press.).
- Cheung, R. K.-M. and W. B. Powell (2000). Shape—a stochastic hybrid approximation procedure for two-stage stochastic programs. *Operations Research* 48(1), 73–79.
- Clark, A. J. and H. Scarf (1960). Optimal policies for a multi-echelon inventory problem. *Management science* 6(4), 475–490.
- De Farias, D. P. and B. Van Roy (2000). On the existence of fixed points for approximate value iteration and temporal-difference learning. *Journal of Optimization theory and Applications* 105(3), 589–608.
- De Farias, D. P. and B. Van Roy (2003). The linear programming approach to approximate dynamic programming. *Operations Research* 51(6), 850–865.
- De Farias, D. P. and B. Van Roy (2004). On constraint sampling in the linear programming approach to approximate dynamic programming. *Mathematics of Operations Research* 29(3), 462–478.
- Desai, V. V., V. F. Farias, and C. C. Moallemi (2012a). Approximate dynamic programming via a smoothed linear program. *Operations Research* 60(3), 655–674.
- Desai, V. V., V. F. Farias, and C. C. Moallemi (2012b). Pathwise optimization for optimal stopping problems. *Management Science* 1909(1996), 1–17.
- Dette, H., S. Hoderlein, and N. Neumeyer (2016). Testing multivariate economic restrictions using quantiles: the example of slutsky negative semidefiniteness. *Journal of Econometrics* 191(1), 129–144.
- Diewert, W. (1974). Applications of duality theory. In M. Intriligator and D. Kendrick (Eds.), *Frontiers of Quantitative Economics, Vol. II*, pp. 106–171. Amsterdam: North-Holland.
- Diewert, W. and T. J. Wales (1987). Flexible functional forms and global curvature conditions. *Econometrica* 55(1), 43–68.
- Diewert, W. E. (1973). Afriat and revealed preference theory. *The Review of Economic Studies* 40(3), 419–425.
- Donohue, C. J. and J. R. Birge (2006). The abridged nested decomposition method for multi-stage stochastic linear programs with relatively complete recourse. *Algorithmic Operations Research* 1(1), 20–30.
- Du, P., C. F. Parmeter, and J. S. Racine (2013). Nonparametric kernel regression with multiple predictors and multiple shape constraints. *Statistica Sinica* 23(3), 1347–1371.
- Dykstra, R. L. (1983). An algorithm for restricted least squares regression. *Journal of the American Statistical Association* 78(384), 837–842.
- Evans, R. V. (1967). Inventory control of a multiproduct system with a limited production resource. *Naval Research Logistics (NRL)* 14(2), 173–184.
- Florens, J., J. Racine, and S. Centorrino (2018). Nonparametric instrumental variable derivative estimation. *Journal of Nonparametric Statistics*, 1–24.
- Førsund, F. R. and L. Hjalmarsson (2004). Are all scales optimal in dea? theory and empirical evidence. *Journal of Productivity Analysis* 21(1), 25–48.
- Frisch, R. (1964). *Theory of production*. Springer Science & Business Media.
- Gallant, A. (1981). On the bias in flexible functional forms and an essentially unbiased form. the fourier flexible form. *Journal of Econometrics* 15(2), 211–245.
- Gallant, A. R. and G. H. Golub (1984). Imposing curvature restrictions on flexible functional

- forms. *Journal of Econometrics* 26(3), 295–321.
- Geramifard, A., T. J. Walsh, S. Tellex, G. Chowdhary, N. Roy, and J. P. How (2013). A tutorial on linear function approximators for dynamic programming and reinforcement learning. *Foundations and Trends in Machine Learning* 6(4), 375–454.
- Godfrey, G. A. and W. B. Powell (2001). An adaptive, distribution-free algorithm for the newsvendor problem with censored demands, with applications to inventory and distribution. *Management Science* 47(8), 1101–1112.
- Goldman, S. and P. Ruud (1993). Nonparametric multivariate regression subject to constraint. Technical Report 93-213, , Department of Economics, University of California at Berkeley, Berkeley, CA.
- Green, P. J. (1995). Reversible jump markov chain monte carlo computation and bayesian model determination. *Biometrika* 82(4), 711–732.
- Griliches, Z. and J. Mairesse (1995). Production functions: the search for identification. Technical report, National Bureau of Economic Research.
- Groeneboom, P., G. Jongbloed, and J. A. Wellner (2001). Estimation of a convex function: characterizations and asymptotic theory. *The Annals of Statistics* 29(6), 1653–1698.
- Guntuboyina, A. and B. Sen (2013). Covering numbers for convex functions. *IEEE Transactions on Information Theory* 59(4), 1957–1965.
- Guntuboyina, A. and B. Sen (2015). Global risk bounds and adaptation in univariate convex regression. *Probability Theory and Related Fields* 163(1-2), 379–411.
- Hall, P. and L.-S. Huang (2001). Nonparametric kernel regression subject to monotonicity constraints. *The Annals of Statistics* 29(3), 624–647.
- Hannah, L. and D. Dunson (2012). Ensemble methods for convex regression with applications to geometric programming based circuit design. arXiv preprint [arXiv:1206.4645](https://arxiv.org/abs/1206.4645).
- Hannah, L. A. and D. B. Dunson (2011). Bayesian nonparametric multivariate convex regression. arXiv preprint [arXiv:1109.0322](https://arxiv.org/abs/1109.0322).
- Hannah, L. A. and D. B. Dunson (2013). Multivariate convex regression with adaptive partitioning. *The Journal of Machine Learning Research* 14(1), 3261–3294.
- Hausman, J. A. and W. K. Newey (2016). Individual heterogeneity and average welfare. *Econometrica* 84(3), 1225–1248.
- Higle, J. L. and S. Sen (1991). Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research* 16(3), 650–669.
- Higle, J. L. and S. Sen (1996). Stopping rules for stochastic decomposition. In *Stochastic Decomposition*, pp. 131–164. Springer.
- Hildreth, C. (1954). Point estimates of ordinates of concave functions. *Journal of the American Statistical Association* 49(267), 598–619.
- Huh, W. T. and P. Rusmevichientong (2009). A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* 34(1), 103–123.
- Hwangbo, H., A. L. Johnson, and Y. Ding (2015). Power curve estimation: Functional estimation imposing the regular ultra passum law. SSRN working paper available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2621033.
- Jiang, D. R. and W. B. Powell (2015). An approximate dynamic programming algorithm for monotone value functions. *Operations Research* 63(6), 1489–1511.
- Juditsky, A. and A. Nemirovski (2002). On nonparametric tests of positivity/monotonicity/convexity. *Annals of statistics* 30(2), 498–527.
- Kamien, M. I. and N. L. Schwartz (1981). *Dynamic optimization*. North-Holland, New York, NY.
- Kuchibhotla, A. K., R. K. Patra, and B. Sen (2017). Efficient estimation in convex single index models. arXiv preprint [arXiv:1708.00145](https://arxiv.org/abs/1708.00145).
- Kunnumkal, S. and H. Topaloglu (2008a). Exploiting the structural properties of the underlying markov decision problem in the q-learning algorithm. *INFORMS Journal on Computing* 20(2), 288–301.
- Kunnumkal, S. and H. Topaloglu (2008b). Using stochastic approximation methods to compute optimal base-stock levels in inventory control problems. *Operations Research* 56(3), 646–664.
- Kuosmanen, T. (2008). Representation theorem for convex nonparametric least squares. *The Econometrics Journal* 11(2), 308–325.
- Kuosmanen, T. and A. Johnson (2017). Modeling joint production of multiple outputs in stoned: Directional distance function approach. *European Journal of Operational Research* 262(2), 792–801.

- Kurt, M. and J. P. Kharoufeh (2010). Monotone optimal replacement policies for a Markovian deteriorating system in a controllable environment. *Operations Research Letters* 38(4), 273–279.
- Kurt, M. and L. M. Maillart (2009). Structured replacement policies for a markov-modulated shock model. *Operations Research Letters* 37(4), 280–284.
- Kushner, H. J. and G. G. Yin (2003). *Stochastic approximation and recursive algorithms and applications*, Volume 35. Springer.
- Lee, C.-Y., A. L. Johnson, E. Moreno-Centeno, and T. Kuosmanen (2013). A more efficient algorithm for convex nonparametric least squares. *European Journal of Operational Research* 227(2), 391–400.
- Lewbel, A. (2001). Demand systems with and without errors. *American Economic Review* 91(3), 611–618.
- Li, Q. and J. S. Racine (2007). *Nonparametric econometrics: theory and practice*. Princeton University Press.
- Lim, E. (2014). On convergence rates of convex regression in multiple dimensions. *INFORMS Journal on Computing* 26(3), 616–628.
- Lim, E. and P. W. Glynn (2012). Consistency of multidimensional convex regression. *Operations Research* 60(1), 196–208.
- Linowsky, K. and A. B. Philpott (2005). On the convergence of sampling-based decomposition algorithms for multistage stochastic programs. *Journal of Optimization Theory and Applications* 125(2), 349–366.
- Löhndorf, N., D. Wozabal, and S. Minner (2013). Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research* 61(4), 810–823.
- Luo, Y. and E. Lim (2016). On consistency of absolute deviations estimators of convex functions. *International Journal of Statistics and Probability* 5(2), 1.
- Maceira, M. E. P., L. G. B. Marzano, D. D. J. Penna, A. L. Diniz, and T. C. Justino (2015). Application of CVaR risk aversion approach in the expansion and operation planning and for setting the spot price in the Brazilian hydrothermal interconnected system. *International Journal of Electrical Power and Energy Systems* 72, 126–135.
- Magnani, A. and S. P. Boyd (2009). Convex piecewise-linear fitting. *Optimization and Engineering* 10(1), 1–17.
- Mak, W.-K., D. P. Morton, and R. K. Wood (1999). Monte carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters* 24(1-2), 47–56.
- Mammen, E. (1991). Nonparametric regression under qualitative smoothness assumptions. *The Annals of Statistics* 19(2), 741–759.
- Matzkin, R. L. (1991). Semiparametric estimation of monotone and concave utility functions for polychotomous choice models. *Econometrica: Journal of the Econometric Society* 59(5), 1315–1327.
- Matzkin, R. L. (1994). Restrictions of economic theory in nonparametric methods. *Handbook of econometrics* 4, 2523–2558.
- Mazumder, R., A. Choudhury, G. Iyengar, and B. Sen (2015). A computational framework for multivariate convex regression and its variants. arXiv preprint [arXiv:1509.08165](https://arxiv.org/abs/1509.08165).
- Michaelides, P. G., E. G. Tsionas, A. T. Vouldis, and K. N. Konstantakis (2015). Global approximation to arbitrary cost functions: A bayesian approach with application to us banking. *European Journal of Operational Research* 241(1), 148–160.
- Mnih, V., K. Kavukcuoglu, D. Silver, A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, and S. Petersen (2015). Human-level control through deep reinforcement learning. *Nature* 518(7540), 529–533.
- Mundlak, Y. (1963). Specification and estimation of multiproduct production functions. *Journal of Farm Economics* 45(2), 433–443.
- Nascimento, J. M. and W. B. Powell (2009). An optimal approximate dynamic programming algorithm for the lagged asset acquisition problem. *Mathematics of Operations Research* 34(1), 210–237.
- Nascimento, J. M. and W. B. Powell (2010). Dynamic programming models and algorithms for the mutual fund cash balance problem. *Management Science* 56(5), 801–815.
- Nascimento, J. M. and W. B. Powell (2013). An optimal approximate dynamic programming algorithm for concave, scalar storage problems with vector-valued controls. *IEEE Transactions on Automatic Control* 58(12), 2995–3010.

- Neave, E. H. (1970). The stochastic cash balance problem with fixed costs for increases and decreases. *Management Science* 16(7), 472–490.
- Ormonoit, D. and Á. Sen (2002). Kernel-based reinforcement learning. *Machine Learning* 49(2–3), 161–178.
- Papadaki, K. P. and W. B. Powell (2002). Exploiting structure in adaptive dynamic programming algorithms for a stochastic batch service problem. *European Journal of Operational Research* 142(1), 108–127.
- Pereira, M. V. F. and L. M. V. G. Pinto (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52, 359–375.
- Perloff, J. M. (2018). *Microeconomics. Eighth edition.* Pearson series in economics. Boston : Pearson, [2015].
- Philpott, A. and Z. Guan (2008). On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters* 36(4), 450–455.
- Philpott, A. B. and V. L. de Matos (2012). Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research* 218(2), 470–483.
- Philpott, A. B., V. L. de Matos, and E. Finardi (2013). On solving multistage stochastic programs with coherent risk measures. *Operations Research* 61(4), 957–970.
- Pierskalla, W. P. and J. A. Voelker (1976). A survey of maintenance models: the control and surveillance of deteriorating systems. *Naval Research Logistics Quarterly* 23(3), 353–388.
- Porteus, E. L. (2002). *Foundations of stochastic inventory theory.* Stanford University Press.
- Powell, W. B. (2011). *Approximate dynamic programming: Solving the curses of dimensionality* (2nd ed.). Wiley.
- Powell, W. B., A. Ruszczyński, and H. Topaloglu (2004). Learning algorithms for separable approximations of discrete stochastic optimization problems. *Mathematics of Operations Research* 29(4), 814–836.
- Pritchard, G., B. Philpott, and J. Neame (2005). Hydroelectric reservoir optimization in a pool market. *Mathematical Programming* 103(3), 445–461.
- Puterman, M. L. (1994). *Markov decision processes: Discrete stochastic dynamic programming.* New York: Wiley.
- Pya, N. and S. N. Wood (2015). Shape constrained additive models. *Statistics and Computing* 25(3), 543–559.
- Robertson, T., F. Wright, and R. Dykstra (1988). *Order Restricted Statistical Inference.* Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society* 55(5), 999–1033.
- Ryan, D. L. and T. J. Wales (2000). Imposing local concavity in the translog and generalized leontief cost functions. *Economics Letters* 67(3), 253–260.
- Samuelson, P. A. (1938). A note on the pure theory of consumer’s behaviour. *Economica* 5(17), 61–71.
- Samuelson, P. A. (1950). The problem of integrability in utility theory. *Economica* 17(68), 355–385.
- Sarith, B. and A. Maindiratta (1997). On the consistency of maximum likelihood estimation of monotone and concave production frontiers. *Journal of Productivity Analysis* 8(3), 239–246.
- Scarf, H. (1960). The optimality of (s, S) policies in the dynamic inventory problem. Technical report, Applied Mathematics and Statistics Laboratory, Stanford University.
- Seijo, E. and B. Sen (2011). Nonparametric least squares estimation of a multivariate convex regression function. *The Annals of Statistics* 39(3), 1633–1657.
- Sen, S. and Z. Zhou (2014). Multistage stochastic decomposition: a bridge between stochastic programming and approximate dynamic programming. *SIAM Journal on Optimization* 24(1), 127–153.
- Shapiro, A. (2011). Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research* 209(1), 63–72.
- Shapiro, A., W. Tekaya, J. P. da Costa, and M. P. Soares (2013). Risk neutral and risk averse stochastic dual dynamic programming method. *European Journal of Operational Research* 224(2), 375–391.
- Silver, D., A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hass-

- abis (2016). Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587), 484–489.
- Stokey, N. L., R. E. Lucas, Jr., and E. C. Prescott (1989). *Recursive methods in economic dynamics*. Harvard University Press.
- Sutton, R. S. and A. G. Barto (1998). *Reinforcement learning: An introduction*. MIT Press.
- Sutton, R. S., D. A. McAllester, S. P. Singh, and Y. Mansour (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pp. 1057–1063.
- Tripathi, G. (2000). Local semiparametric efficiency bounds under shape restrictions. *Econometric Theory* 16(5), 729–739.
- Tsitsiklis, J. N. and B. V. Roy (1996). Feature-based methods for large scale dynamic programming. *Machine Learning* 22, 59–94.
- Tsitsiklis, J. N. and B. Van Roy (1999). Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control* 44(October), 1840–1851.
- Varian, H. R. (1982). The nonparametric approach to demand analysis. *Econometrica* 50(4), 945–973.
- Varian, H. R. (1984). The nonparametric approach to production analysis. *Econometrica* 52(3), 579–597.
- Varian, H. R. (1992). *Microeconomic Analysis*. WW Norton&Company Inc.
- Villalobos, M. and G. Wahba (1987). Inequality-constrained multivariate smoothing splines with application to the estimation of posterior probabilities. *Journal of the American Statistical Association* 82(397), 239–248.
- Watkins, C. J. and P. Dayan (1992). Q-learning. *Machine Learning* 8(3-4), 279–292.
- Wu, C.-F. (1982). Some algorithms for concave and isotonic regression. *Studies in the Management Sciences* 19, 105–116.
- Wu, X. and R. Sickles (2017). Semiparametric estimation under shape constraints. *Econometrics and Statistics*.
- Yagi, D., Y. Chen, A. L. Johnson, and T. Kuosmanen (2018). Shape constrained kernel-weighted least squares: Application to production function estimation for chilean manufacturing industries. *Accepted at the Journal of Business and Economic Statistics*, arXiv working paper available at: <https://arxiv.org/abs/1604.06003>.
- Yagi, D., Y. Chen, A. L. Johnson, and H. Morita (2018). Iterative nonparametric s-shape estimation. arXiv preprint.
- Yatchew, A. (2003). *Semiparametric regression for the applied econometrician*. Cambridge University Press.
- Zhang, H., X. Chao, and C. Shi (2017). Perishable inventory problems: Convexity results for base-stock policies and learning algorithms under censored demand. *Operations Research*.